

Statistical Physics of Information Measures

Neri Merhav

Department of Electrical Engineering

Technion – Israel Institute of Technology

Technion City, Haifa, Israel

Partly joint work with **D. Guo** (Northwestern U.) and **S. Shamai** (Technion).

Physics of Algorithms '09, Santa Fe, NM, USA, Aug. 31 – Sep. 4, 2009

Outline

Relations between Information Theory (IT) and statistical physics:

- **Conceptual aspects** – relations between principles in the two areas.
- **Technical aspects** – identifying similar mathematical formalisms and borrowing techniques.

In this talk we:

- Briefly review basic background in IT.
- Discuss some **physics of the Shannon limits**.
- Briefly review basic background in estimation theory.
- Touch upon **statistical physics** of **signal estimation** via the mutual information.

First Part:

Physics of the Shannon Limits

The Shannon Limits

- Lossless data compression:

compression ratio $\geq H$ = entropy.

- Lossy compression:

compression ratio $\geq R(D)$ = rate–distortion func.

- Channel coding:

coding rate $\leq C$ = channel capacity.

- Joint source–channel coding:

decoding error $\geq R^{-1}(C)$ = distortion–rate func. at rate C .

- etc. etc. etc.

The Information Inequality

Each of the above–mentioned **fundamental limits** of IT, as well as many others, is based on the **information inequality** in some form:

For any two distributions, P and Q , over an alphabet \mathcal{X} :

$$D(P\|Q) \triangleq \sum_x P(x) \log \frac{P(x)}{Q(x)} \geq 0.$$

In physics, it is known as the **Gibbs inequality**.

The Gibbs Inequality

Let $\mathcal{E}_0(x)$ and $\mathcal{E}_1(x)$ be two Hamiltonians of a system. For a given β , let

$$P_i(x) = \frac{e^{-\beta \mathcal{E}_i(x)}}{Z_i}, \quad Z_i = \sum_x e^{-\beta \mathcal{E}_i(x)}, \quad i = 0, 1.$$

Then,

$$\begin{aligned} 0 &\leq D(P_0 \| P_1) = \left\langle \ln \frac{e^{-\beta \mathcal{E}_0(X)} / Z_0}{e^{-\beta \mathcal{E}_1(X)} / Z_1} \right\rangle_0 \\ &= \ln Z_1 - \ln Z_0 + \beta \langle \mathcal{E}_1(X) - \mathcal{E}_0(X) \rangle_0 \end{aligned}$$

or

$$\begin{aligned} \langle \mathcal{E}_1(X) - \mathcal{E}_0(X) \rangle_0 &\geq kT \ln Z_0 - kT \ln Z_1 \\ &= F_1 - F_0 \end{aligned}$$

Interpretation of $\langle \mathcal{E}_1(X) - \mathcal{E}_0(X) \rangle_0 \geq \Delta F$

- A system with Hamiltonian $\mathcal{E}_0(x)$ – in equilibrium $\forall t < 0$.
Free energy $= -kT \ln Z_0$.
- At $t = 0$, the Hamiltonian **jumps**, by $W = \mathcal{E}_1(x) - \mathcal{E}_0(x)$: from $\mathcal{E}_0(x)$ to $\mathcal{E}_1(x)$ – by **abruptly** applying a **force**. Energy injected:
 $\langle W \rangle_0 = \langle \mathcal{E}_1(X) - \mathcal{E}_0(X) \rangle_0$.
- New system, with Hamiltonian \mathcal{E}_1 , equilibrates.
Free energy $= -kT \ln Z_1$.

Gibbs inequality: $\langle W \rangle_0 \geq \Delta F$.

$$\langle W \rangle_0 - \Delta F = kT \cdot D(P_0 \| P_1)$$

is the **dissipated energy** = entropy production (system + environment) due to **irreversibility** of the **abruptly** applied force.

Example – Data Compression and the Ising Model

Let $\mathbf{X} \in \{-1, +1\}^n \sim$ Markov chain $P_0(\mathbf{x}) = \prod_i P_0(x_i|x_{i-1})$ with

$$P_0(x|x') = \frac{\exp(Jx \cdot x')}{Z_0}, \quad x, x' \in \{-1, +1\}$$

Code designer thinks that $\mathbf{X} \sim$ Markov with parameters:

$$P_1(x|x') = \frac{\exp(Jx \cdot x' + Kx)}{Z_1(x')}.$$

$D(P_0\|P_1) =$ **loss** in compression due to **mismatch**. Easy to see that

$$\mathcal{E}_0(\mathbf{x}) = -J \cdot \sum_i x_i x_{i-1}; \quad \mathcal{E}_1(\mathbf{x}) = -J \cdot \sum_i x_i x_{i-1} - B \cdot \sum_i x_i$$

where

$$B = K + \frac{1}{2} \ln \frac{\cosh(J - K)}{\cosh(J + K)}.$$

Thus, $W = -B \cdot \sum_i x_i$ means an abrupt application of the magnetic field B .

Physics of the Data Processing Theorem (DPT)

Mutual information: Let $(U, V) \sim P(u, v)$:

$$I(U; V) \equiv \left\langle \log \frac{P(U, V)}{P(U)P(V)} \right\rangle.$$

DPT:

$$X \rightarrow U \rightarrow V \text{ Markov chain} \implies I(X; U) \geq I(X; V).$$

Pf:

$$I(X; U) - I(X; V) = \left\langle D(P_{X|U,V}(\cdot|U, V) \| P_{X|V}(\cdot|V)) \right\rangle \geq 0. \quad \square$$

Supports most, if not \forall , Shannon limits.

Physics of the DPT (Cont'd)

Let $\beta = 1$. Given (u, v) , let

$$\mathcal{E}_0(x) = -\ln P(x|u, v) = -\ln P(x|u); \quad \mathcal{E}_1(x) = -\ln P(x|v).$$

$$Z_0 = \sum_x e^{-1 \cdot [-\ln P(x|u, v)]} = \sum_x P(x|u, v) = 1$$

and similarly, $Z_1 = 1$. Thus, $F_0 = F_1 = 0$, and so, $\Delta F = 0$.

After averaging over P_{UV} :

$$\begin{aligned} \langle W(X) \rangle_0 &= \langle -\ln P(X|V) + \ln P(X|U) \rangle \\ &= H(X|V) - H(X|U) \\ &= I(X; U) - I(X; V). \end{aligned}$$

$$\langle W \rangle_0 = I(X; U) - I(X; V) \geq 0 = \Delta F.$$

Discussion

The relation

$$\langle W \rangle_0 - \Delta F = kT \cdot D(P_0 \| P_1) \geq 0$$

is known (Jarzynski '97, Crooks '99, ..., Kawai *et. al.* '07), but with different physical interpretations, which require some limitations.

Present interpretation – holds generally; Applied in particular to the DPT.

In our case:

- Maximum irreversibility: $\langle W \rangle_0$ – fully dissipated: $\Delta F = 0$.
- All dissipation – in the system, none in the environment:

$$\langle W \rangle_0 = T\Delta S = 1 \cdot [H(X|V) - H(X|U)].$$

- Rate loss due to gap between mutual informations:
irreversible process \iff irreversible info: $I(X;U) > I(X;V) \longrightarrow U$
cannot be retrieved from V .

Relation to Jarzynski's Equality

Let

$$\mathcal{E}_\lambda(x) = \mathcal{E}_0(x) + \lambda[\mathcal{E}_1(x) - \mathcal{E}_0(x)]$$

interpolate \mathcal{E}_0 and \mathcal{E}_1 . λ – a generalized **force**.

Jarzynski's equality (1997): \forall protocol $\{\lambda_t\}$ with $\lambda_t = 0 \forall t \leq 0$ and $\lambda_t = 1 \forall t \geq \tau$ ($\tau \geq 0$), the injected energy

$$W = \int_0^\tau d\lambda_t [\mathcal{E}_1(x_t) - \mathcal{E}_0(x_t)]$$

satisfies

$$\langle e^{-\beta W} \rangle = e^{-\beta \Delta F}.$$

Jensen: $\langle e^{-\beta W} \rangle \geq \exp\{-\beta \langle W \rangle\}$ so, $\langle W \rangle \geq \Delta F$ more generally.

Equality – for a **reversible** process – W = deterministic.

Informational Jarzynski Equality

Taking

$$\mathcal{E}_0(x) = -\ln P_0(x), \quad \mathcal{E}_1(x) = -\ln P_1(x), \quad \beta = 1$$

and defining a “protocol” $0 \equiv \lambda_0 \rightarrow \lambda_1 \rightarrow \dots \rightarrow \lambda_n \equiv 1$, and

$$W = \sum_{i=0}^{n-1} (\lambda_{i+1} - \lambda_i) \ln \frac{P_0(X_i)}{P_1(X_i)}, \quad X_i \sim P_{\lambda_i} \propto P_0^{1-\lambda_i} P_1^{\lambda_i},$$

one can show:

$$\langle e^{-W} \rangle = 1 = e^{-\Delta F}.$$

Jensen: [generalized information inequality](#):

$$\int_0^1 d\lambda_t \left\langle \ln \frac{P_0(X)}{P_1(X)} \right\rangle_{\lambda_t} \geq 0.$$

Summary of First Part

- Suboptimum commun. system \Longleftrightarrow irreversible process.
- Info rate loss \Longleftrightarrow dissipated energy \rightarrow entropy \uparrow
- Fundamental limits of IT \Longleftrightarrow second law.
- Possible implications of Jarzynski's equality in IT.

Second Part:

Statistical Physics of Signal Estimation via the Mutual Information

Signal Estimation – Preliminaries

Let

$$\mathbf{Y} = \mathbf{X} + \mathbf{Z} \quad (\text{all vectors in } \mathbb{R}^n)$$

where \mathbf{X} is the **desired signal** and \mathbf{Z} is **noise** $\perp \mathbf{X}$.

Estimator: any function $\hat{\mathbf{X}} = f(\mathbf{Y})$. We want $\hat{\mathbf{X}}$ as ‘close’ as possible to \mathbf{X} .

$$\text{mean square error} = \langle \|\mathbf{X} - \hat{\mathbf{X}}\|^2 \rangle = \langle \|\mathbf{X} - f(\mathbf{Y})\|^2 \rangle.$$

A fundamental result: **minimum mean square error** (MMSE) = **conditional mean**:

$$\mathbf{X}^* = f^*(\mathbf{y}) = \langle \mathbf{X} \rangle_{\mathbf{Y}=\mathbf{y}} \equiv \int d\mathbf{x} \cdot \mathbf{x} P(\mathbf{x}|\mathbf{y}).$$

Normally – difficult to apply \mathbf{X}^* and assess performance.

\mathbf{X}^* and MMSE may exhibit irregularities – **threshold effects** \longleftrightarrow **phase transitions** in analogous physical systems. Motivates a **statistical–mechanical perspective**.

The I-MMSE Relation

[Guo–Shamai–Verdú 2005]: for $\mathbf{Y} = \mathbf{X} + \mathbf{Z}$, $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I} \cdot 1/\beta)$, regardless of $P(\mathbf{X})$:

$$\text{mmse}(\mathbf{X}|\mathbf{Y}) = 2 \cdot \frac{\text{d}}{\text{d}\beta} I(\mathbf{X}; \mathbf{Y}),$$

where $\text{mmse}(\mathbf{X}|\mathbf{Y}) \equiv \langle \|\mathbf{X} - f^*(\mathbf{Y})\|^2 \rangle$.

Simple example: If $\mathbf{X} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$,

$$\frac{I(\mathbf{X}; \mathbf{Y})}{n} = \frac{1}{2} \log(1 + \beta \sigma^2)$$

$$\implies \frac{\text{mmse}(\mathbf{X}|\mathbf{Y})}{n} = \frac{\sigma^2}{1 + \beta \sigma^2}.$$

MMSE – calculated using stat-mech via the mutual info and I-MMSE relation

\implies

Statistical Physics of the MMSE

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}) &= \left\langle \log \frac{P(\mathbf{X}|\mathbf{Y})}{P(\mathbf{X})} \right\rangle_{\beta} \\ &= \left\langle \log \frac{\exp\{-\beta \|\mathbf{Y} - \mathbf{X}\|^2/2\}}{\sum_{\mathbf{x}} P(\mathbf{x}) \exp\{-\beta \|\mathbf{Y} - \mathbf{x}\|^2/2\}} \right\rangle_{\beta} \\ &= -\frac{n}{2} - \langle \log Z(\beta|\mathbf{Y}) \rangle_{\beta} \end{aligned}$$

where

$$Z(\beta|\mathbf{Y}) = \sum_{\mathbf{x}} P(\mathbf{x}) \exp\{-\beta \|\mathbf{Y} - \mathbf{x}\|^2/2\},$$

and so,

$$\text{mmse}(\mathbf{X}|\mathbf{Y}) = 2 \cdot \frac{\text{d}I(\mathbf{X}; \mathbf{Y})}{\text{d}\beta} = -2 \frac{\partial}{\partial \beta} \langle \log Z(\beta|\mathbf{Y}) \rangle_{\beta}.$$

Similar to [internal energy](#), but here also $\langle \cdot \rangle_{\beta}$ depends on β .

Statistical Physics of the MMSE (Cont'd)

A more detailed derivation yields:

$$\text{mmse}(\mathbf{X}|\mathbf{Y}) = \frac{n}{\beta} + \text{Cov}\{\|\mathbf{Y} - \mathbf{X}\|^2, \log Z(\beta|\mathbf{Y})\}$$

- The term $n/\beta \sim$ energy equipartition theorem.
- Covariance term – dependence of $\langle \cdot \rangle_\beta$ on β .

Statistical Physics of the MMSE (Cont'd)

$$\begin{aligned}\text{In stat. mech: } \Sigma(\beta) &= \log Z(\beta) + \beta \langle \mathcal{E}(X) \rangle \\ &= \log Z(\beta) - \beta \frac{d \log Z(\beta)}{d\beta} \quad \Longleftarrow \text{diff. eq.}\end{aligned}$$

$$\log Z(\beta) = -\beta E_0 + \beta \cdot \int_{\beta}^{\infty} \frac{d\hat{\beta} \cdot \Sigma(\hat{\beta})}{\hat{\beta}^2}; \quad E_0 = \text{ground-state energy}$$

$$\Rightarrow E = -\frac{d \log Z(\beta)}{d\beta} = \left[E_0 - \int_{\beta}^{\infty} \frac{d\hat{\beta} \cdot \Sigma(\hat{\beta})}{\hat{\beta}^2} \right] + \frac{\Sigma(\beta)}{\beta}$$

Similarly for $\langle \log Z(\beta | \mathbf{Y}) \rangle_{\beta}$ except that

$$\Sigma(\beta) \Longleftarrow \frac{\beta}{2} \mathbf{Cov}\{\|\mathbf{Y} - \mathbf{X}\|^2, \log Z(\beta | \mathbf{Y})\} - I(\mathbf{X}; \mathbf{Y})$$

$$E_0 \Longleftarrow \frac{1}{2} \left\langle \min_{\mathbf{x}} \|\mathbf{Y} - \mathbf{x}\|^2 \right\rangle_{\beta}.$$

Examples

Example 1 – Random Codebook on a Sphere Surface

$$\mathbf{Y} = \mathbf{X} + \mathbf{Z}; \quad \mathbf{X} \sim \text{Unif}\{\mathbf{x}_1, \dots, \mathbf{x}_M\}, \quad M = e^{nR}$$

Codewords: randomly drawn independently uniformly on $\text{Surf}(\sqrt{n\sigma^2})$.

$$\lim_{n \rightarrow \infty} \frac{\langle I(\mathbf{X}; \mathbf{Y}) \rangle}{n} = \begin{cases} \frac{1}{2} \log(1 + \beta\sigma^2) & \beta < \beta_R \\ R & \beta \geq \beta_R \end{cases}$$

where β_R is the solution to the eqn $R = \frac{1}{2} \log(1 + \beta\sigma^2)$. Thus,

$$\lim_{n \rightarrow \infty} \frac{\text{mmse}(\mathbf{X}|\mathbf{Y})}{n} = \begin{cases} \frac{\sigma^2}{1 + \beta\sigma^2} & \beta < \beta_R \\ 0 & \beta \geq \beta_R \end{cases}$$

A 1st-order ϕ transition in MMSE: At high temp. behaves as if \mathbf{X} was Gaussian and at $\beta = \beta_R$ jumps to zero!

Examples (Cont'd)

Example 2 – Sparse Signals

$$X_i = \left(\frac{1 - \mu_i}{2} \right) U_i, \quad i = 1, \dots, n$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n) \sim P(\boldsymbol{\mu})$ are **binary** $\{\pm 1\}$; $U_i \sim \mathcal{N}(0, \sigma^2)$ – i.i.d. $\perp \boldsymbol{\mu}$.

$$\begin{aligned} Z(\beta|\mathbf{y}) &= \int_{\mathbb{R}^n} d\mathbf{x} P(\mathbf{x}) \exp\{-\beta \|\mathbf{y} - \mathbf{x}\|^2/2\} \iff P(\mathbf{x}) = \sum_{\boldsymbol{\mu}} P(\boldsymbol{\mu}) P(\mathbf{x}|\boldsymbol{\mu}) \\ &= \sum_{\boldsymbol{\mu}} P(\boldsymbol{\mu}) \exp\left\{-\frac{1}{2} \sum_{i=1}^n \text{func}(y_i, \mu_i, q)\right\} \iff q \equiv \beta\sigma^2 \\ &= \text{const.} \times \sum_{\boldsymbol{\mu}} P(\boldsymbol{\mu}) \cdot \exp\left\{\sum_{i=1}^n \mu_i h_i\right\} \quad h_i = \text{func}(y_i) \end{aligned}$$

Sum over $\{\boldsymbol{\mu}\} \equiv \hat{Z}(\beta|\mathbf{y})$: “partition function” of spins in a **random field** $\{h_i\}$.

Example 2 (Cont'd)

Let $P(\boldsymbol{\mu}) \propto \exp\{nf[m(\boldsymbol{\mu})]\}$ where $m(\boldsymbol{\mu}) \equiv \frac{1}{n} \sum_i \mu_i$ and $f[m]$ is 'nice'.

$$\hat{Z}(\beta|\mathbf{y}) \propto \sum_{\boldsymbol{\mu}} \exp \left\{ n \left[f[m(\boldsymbol{\mu})] + \frac{1}{n} \sum_i \mu_i h_i \right] \right\}$$

\hat{Z} is dominated by configurations with magnetization m^* , solving the zero-derivative equation

$$m = \langle \tanh(f'[m] + H) \rangle$$

where H is a RV pertaining to h_i . m^* = local maximum if:

$$\langle \tanh^2(f'[m^*] + H) \rangle > 1 - \frac{1}{f''[m^*]}.$$

When this becomes equality (and then reversed), m^* ceases to dominate \hat{Z} (critical point) \implies dominant magnetization jumps elsewhere.

Example 2 (Cont'd)

Consider the case

$$f[m] = am + \frac{bm^2}{2}$$

\hat{Z} – similar to the [random-field Curie-Weiss](#) (RFCW) model.

We analyze the mutual info using stat-mech methods, and then derive the MMSE using the I-MMSE relation:

MMSE for Example 2

$$\begin{aligned}
 \overline{\text{mmse}} = & \frac{\sigma^2 q}{2(1+q)^2} + \frac{(1-m_a)\sigma^2}{2} \left[1 - \frac{q(1+q/2)}{(1+q)^2} \right] + \\
 & \frac{1+m_a}{2} \left[\text{Cov}_0\{Y^2, \log[2 \cosh(b\mathbf{m}^* + a + H)]\} + \right. \\
 & \left. \langle H' \tanh(b\mathbf{m}^* + a + H) \rangle_0 \right] + \\
 & + \frac{1-m_a}{2} \left[\frac{1}{(1+q)^2} \cdot \text{Cov}_1\{Y^2, \log[2 \cosh(b\mathbf{m}^* + a + H)]\} + \right. \\
 & \left. \langle H' \tanh(b\mathbf{m}^* + a + H) \rangle_1 \right]
 \end{aligned}$$

where $\langle \cdot \rangle_s$ and Cov_s are w.r.t. $Y \sim \mathcal{N}(0, \sigma^2 s + 1/\beta)$, $s = 0, 1$, and

$$H' = -\frac{\sigma^2}{2(1+q)} + \frac{q(q+2)Y^2}{2(1+q)^2}.$$

Example 2: Discussion

- MMSE depends on m^* : jumps of m^* yield discontinuities in MMSE.
- As m^* jumps, the response of $X^*(Y)$ jumps as well.
- In the C–W model: 1st order transition w.r.t. mag. field and 2nd order transition w.r.t. β . Here – a 1st order transition w.r.t. β because dependence on β is via the “magnetic fields” $\{h_i\}$..
- $b = 0$: i.i.d. spins \implies no ϕ transitions \implies sparsity alone does not cause ϕ transitions.

Conclusion of Second Part

- MMSE calculated using stat. mech. via the mutual info.
- Statistical–mech techniques can be used to inspect **inherent** irregularities in the estimation error, via **phase transitions**.
- Possible to handle situations of mismatch between **true** prior P and **assumed** prior Q .